In the Specification

Page 1 line 26 to page 2 line 13:

STC exploits both the temporal and spatial dimensions for the construction of coding designs which effectively mitigate fading (for improved power efficiency) and are able to ~~capitalise~~ capitalize upon parallel transmission paths within the propagation channel (for improved bandwidth efficiency). The use of multiple propagation paths is fundamental to the concept, and provides the means for better Signal to Noise Ratio (SNR) gain (i.e. power efficiency) and better diversity performance across the radio link. Furthermore, in a rich multipath environment antenna processing can provide access to independent, parallel transmission paths, which are also referred to as 'spatial modes', 'spatial channels' or 'data pipes'.

A2

In Figure 3 there is shown a plot for the throughput (as defined by Foschini and Gans "On Limits of Wireless Communications in a Fading Environment when using Multiple Antennas"', Wireless Personal Communications, Vol. 6, 1998, Kluwer Academic Publishers, pp. 311-335) which is exceeded for 90% of frames (known as the '10% Outage Capacity') from a number of transmit elements (N_T) and a number of receive elements (N_R) for the cases of (N_T,N_R) = (1,1), (2, 2), (4,4) and (1,4) over 10,000 random channel ~~realisations~~realizations for different overall mean SNRs. In order to calculate fundamental per-link capacity bounds, coded data frames of length approaching infinity are assumed, with random, independent, Rayleigh channel ~~realisations~~realizations for each transmitter-receiver path. Thus an (N_R,N_T) matrix of complex channel gains, H, can be constructed with independent complex Gaussian elements. A transmitter would have no prior knowledge of the instantaneous channel ~~realisations~~realizations, and so would send equal-power uncorrelated noise-like waveforms to the different antennas. It can be shown that the fundamental capacity (i.e. maximum throughput for vanishingly small error probability) for such a system would vary from frame-to-frame, as a function of the eigenvalues of the matrix product HH^H.

Page 2 line 24 to page 3 line 3:

It is commonly accepted that the capacity of CDMA cellular systems is inversely proportional to the working E_b/N_o of the subscribers. In the reverse-link case, it is the received E_b which must be ~~minimised~~minimized, whereas for the forward-link our aim is to ~~minimise~~minimize the transmitted E_b (where E_b denotes energy per information bit). Thus the per-link throughput increases as a logarithmic function of link E_b (in the manner prescribed by Shannon, 'Communication in the Presence of Noise' Proc. IRE, vol. 37, pp. 10-21, Jan. 1949), whereas the number of links which can be supported is ~~in~~ inversely proportional to E_b. If we consider the overall network capacity to be the product of the per-link capacity and number of links, then we obtain the characteristic shown in Figure 4. This shows that systems with larger numbers of links, each working at a lower E_b/N_o, and hence lower throughput, will have a larger network capacity (i.e. higher spectral efficiency) than systems working at a higher link E_b/N_o, but with fewer links. This is independent of any statistical multiplexing benefits which apply for large numbers of users, and this further ~~favours~~ favors systems with large numbers of low-SNR/low-throughput links. This analysis relies upon the assumption that the system is not bandwidth limited (i.e. in the case of the CDMA forward link this requires that there be a plentiful supply of Walsh codes).

A3

Page 3 line 30 to page 5 line 10:

Several key implementation issues influence the achievable performance gain using STC. Since STC ~~capitalises~~ capitalizes on the inherent parallelism in the spatial channel, one possible risk concerns the substantial increase in DSP load needed to estimate the multiplicity of channel impulse responses at the terminal.

A4 cont

The fundamental principle of STC, is illustrated in Figure 1._Coded data symbols are transmitted simultaneously from multiple antennas. This simultaneous transmission of modulation symbols from different antennas is termed a Space-Time Symbol (STS). Since modulation symbols can conventionally be represented by a complex

number, an STS can be represented by a vector of complex numbers, with the number of complex elements in the vector equal to the number of transmit antennas. In a Frequency Division Multiple Access/Time Division Multiple Access (FDMA/TDMA) system these simultaneous symbols use the same carrier frequency and same symboling waveform. In a CDMA system—an identical symboling waveform, in the form of an identical spreading code (i.e. Walsh code) is also used.

Assuming a non-dispersive channel, the receiver simultaneously detects all of the elements of a transmitted STS using a single symbol-matched filter per receiver antenna. These detection outputs are built up into a (vector) detection statistic. Thus Figure -1 shows that every element of this vectorial receiver detection statistic is a superposition of the multiple simultaneous transmissions, as seen at each receiver antenna. This cross-coupling between antenna transmissions can be thought of as a form of intersymbol interference (ISI). Each element, viewed over many STS, will thus resemble a somewhat ~~disorganised~~ disorganized Quadrature Amplitude Modulated (QAM) type constellation (in the noise-free case), where the exact form of this constellation depends on the STC encoder structure, STC modulation alphabet and instantaneous channel.

A2 cont

In the case of multiple receive antennas the transmissions from individual transmit antennas, however, may not be totally (spatially) separated at the receiver. This spatial-ISI can be overcome in a number of ways, including; by a form of antenna spatial 'nulling' (analogous to linear 'zero-forcing' time-domain ~~equalisation~~ equalization) or spatial Minimum Mean-Squared Error (MMSE) beamforming (analogous to linear MMSE time-domain ~~equalisation~~ equalization); or, subtraction (analogous to Decision Feedback time-domain ~~equalisation~~ equalization); or, maximum-likelihood sequence estimation (MLSE) of coded Space-Time Symbols (and hence information bits) from all transmit antennas, based on observations of the vectorial receiver decision statistic over the whole coded frame.

Using conventional methods (and with a bounded receiver complexity), waveform design is inevitably a compromise between bandwidth efficiency and power efficiency. However, apparently good STC systems can seemingly 'side-step' the above barrier (Shannon bound) by achieving increased bandwidth efficiency with little or no reduction in power efficiency. In principle, STC transmissions could be successfully received in a link with only a single receiver antenna. However, the power efficiency would be very poor, because it would not be possible to ~~capitalise~~ capitalize on the additional SNR and diversity gain provided by a multiple-antenna terminal design, and due to the lack of multiple parallel decoupled 'data pipes'. Put more simply, and referring again to Figure 1 (assuming only a single receiver antenna), the receiver detection statistic would only have a single element, and it would resemble an extremely high-order QAM constellation (but which changes both in size and shape as the channel fades at the Doppler rate). This high-order constellation would have a large number of constellation points, all lying in the same 1-dimensional complex space. Hence, like QAM, the power efficiency of such a link would be rather poor. In the limit, two constellation points may even lie on top of one another, and the receiver would require redundancy (memory) in the coding in order to separate them.

A4 cont'd

Page 6 lines 6 to 11:

Accurate channel estimation (i.e. estimation of the complex gains between each transmit/receive antenna pair) is important for reliable decoding of space-time codes. This is because the maximum likelihood branch metric *M*(**x**,**y**) for a received vector **y** and ~~hypothesised~~hypothesized transmit symbol **x** is given by

A5

$$M(x,y) = \|y - Hx\|_2^2$$

Page 6 lines 18 to 25:

The maximum likelihood sequence estimator (implemented by the Viterbi algorithm) outputs the sequence that ~~minimises~~minimizes the sum of such branch metrics over the window of interest. The channel parameters are required in calculation of these maximum likelihood

A6 cont

metrics. If these parameters are not known exactly, the discrepancy between the parameters used and the actual channel parameters has an effect similar to having an increased amount of thermal noise, which degrades performance. This is in contrast to a single antenna additive noise channel, in which knowledge of the signal amplitude or noise variance is not required for the Viterbi algorithm (although it is required for the maximum- a posteriori [MAP] algorithm).

A6 cont'd

Page 7 lines 1-13:

A7

Since the channel gains H are not known a-priori, they must be estimated. The conventional method is to transmit a known sequence of data called a training sequence, which is used by the receiver to produce estimates of the channel gains. Since the performance of the decoding step depends on the accuracy of these channel estimates, a minimum mean square error estimator (MMSE) may be used. This would ~~minimise~~minimize the variance of the error on the channel estimates, which in turn ~~minimises~~minimizes the implementation loss described in Equation 1.

The typical ~~behaviour~~behavior of MMSE estimators is that their error variance decreases ~~as~~with the inverse of the length of the training sequence. In other words, approximately twice as many training symbols are needed in order to halve the mean squared error in the channel estimates. Thus the length of the training sequences has an impact on achievable system performance.

Page 12 lines 7-10:

A8

In order that the present invention is more fully understood and to show how the same may be carried into effect, reference shall now be made, by way of example only, to the figures ~~and tables~~ as shown in the accompanying drawing sheets, in which:-

Page 12 lines 24-32:

A9 cont

Figure 13–A and 13B are a flowchart showing the joint channel estimating and decoding technique;

Figure 14 is a plotted simulation result indicating the relationship between pilot power and decoder performance for a conventional receiver;

Figure 15 shows indicative performance results for a decoder in accordance with the invention;

A9 cont'd

~~Table 1~~Figure 16 shows possible code designs for MC-STCC-N for different N, where $N=2^{k-1}$;

~~Table 2~~Figure 17 shows simulated performance of STC algorithms; and

~~Table 3~~Figure 18 shows STC algorithm complexity (at a data rate of 153.6kbps).

Page 17 lines 26-35:

A10

The binary convolutional encoder 900 is shown in greater detail in Figure 7. Incoming binary data 902 grouped into units of 4 information bits, is input into a shift register 904 having a constraint length K, of 9. The rate is one half with two modulo 2 summers, 906, 908 generating two binary coded symbols for each data bit input to the encoder. The code symbols are output so that the code symbol (c_0) encoded with generator function s_0, is output first, and the code symbol (c_1) encoded with generator function s_1, is output last. The state of the convolutional encoder, upon ~~initialisation~~ initialization, shall be the all-zero state. The first code symbol output after ~~initialisation~~ initialization shall be a code symbol encoded with generator function s_0. The generator functions of the code shall be s_0 equals 753 (octal) and s_1 equals 561 (octal) as per ~~Table 1~~Figure 16 (Source Odenwalder (1970) and Larsen (1973)-).

Page 21 lines 6-33:

A11 cont

Step 202; Calculate the initial channel estimate $\hat{H}^{(0)}$. -This may be estimated in conventional ways, for example, based on a known training sequence fed to the transmit antennas- at the beginning of each frame, or based on pilot orthogonal Walsh codes. The estimation may be done using any standard channel estimation algorithm, e.g. least squares estimator. -A further alternative is to add known data into the input data stream which causes the encoder to generate known channel

symbols. This is similar to the training sequence technique but has advantages in terms of reduced complexity.

It should be noted that this 'bootstrap' channel estimate does not have to be very good, since it only has to be good enough to allow the iteration described below, to get started. Thus very few training/pilot symbols are required; -certainly less than in an equivalently specified, conventional receiver.

A11 Cont'd

Step 204; Decode data preferably using a 'soft' algorithm such as the a-posteriori probability (APP) algorithm and $\hat{H}^{(i)}$ current channel estimate to produce a sequence of coded symbol probabilities $\Pr^{(i)}(x_j[l])$ for each transmit antenna j = 1,2,...,t and each symbol interval l = 1,2,...,L. From these probabilities calculate the sequence of coded symbol averages to $\mu_j^{(i)}[l] = \sum_{x \in \aleph} x \Pr^{(i)}(x_j[l] = x)$ produce symbol estimates (step 206). Optionally, 'harden' the probabilities by using a 'slicing' technique, for example, by setting the symbols with the highest probability to 1 and all other symbols to 0 and using these to perform a 'hard' cancellation in step 214 below. Alternatively, the decoding could be performed using a Viterbi decoder or other decoder which directly produces 'hard' outputs.

Page 25 line 11 to Page 26 line 5:

The main candidate STC algorithms have been qualitatively described. The respective performances in quantitative terms, based on simulation results will now be discussed. The simulations were carried out at the symbol-level, for a simplified system, whereby frames of space-time symbols were transmitted over a quasi-static flat Rayleigh fading MIMO channel. Each frame carried 400 information bits (800 in the case of BLAST). Performance is quantified in terms of required E_b/N_o at each receiver antenna for a mean frame-error-rate (FER) of 10%, where E_b is the energy per information bit. There is no simulation of fast power control, so received E_b is also equal to transmitted E_b (making the ~~normalising~~normalizing assumption, that each channel coefficient has a mean square value of unity). For simplicity, it is also

A12 Cont

assumed that there is zero correlation between the various transmit-receive antenna paths, which assumes a certain minimum spacing between antenna elements at both the BTS and mobile. The results of these simulations are shown in ~~Table 2~~Figure 17.

The results of ~~Table 2~~Figure 17 are also plotted in Figure 10, showing the trade-off of E_b/N_o performance (power efficiency) -against spectral efficiency. The missing 'third dimension' on the figure is that of decoder complexity. After conversion from E_b/N_o to E_S/N_o (which is the corresponding SNR metric) the performance results can also be compared with the theoretical performance curves shown earlier. However, it should be remembered that the underlying assumptions are different in the two cases, since, in the case of the fundamental performance curves, the frames are assumed to approach infinite length, with a vanishingly small error rate, and a Space-Time Code which adapts from frame-to-frame, whilst, for the simulation case, there are finite frames, a fixed code, and an allowed non-zero probability of frame error.

A13 cont'd

Examining the results of ~~Table 2~~Figure 17 (and Figure 10) it can be seen that the benchmark has a poor performance in the 1:1 case, with a required E_b/N_o of 11.0dB. This is 8.8dB higher than the static channel performance, and is due to the high fade margin required due to the lack of diversity. This performance can be improved either by increasing the number of receiver antennas (yielding both SNR gain and diversity gain) or by increasing the number of transmit antennas and using STTD, or both. In the latter case, for STTD 2:2, we see a required E_b/N_o of only 2.2dB.

Page 26 line 27 to page 27 line 32:

The uncoded BLAST scheme shown in ~~Table 2~~Figure 17 performs poorly in terms of E_b/N_o, but has a high spectral efficiency due to the absence of redundancy, and the use of parallel transmission paths. Separate frames of data (each carrying 400 information bits) are transmitted from each of the two transmit antennas. The stronger of the two frames is detected first, by spatially nulling out (zero forcing) the

A13 cont

weaker transmission, before detecting the bits. A 'genie-aided' subtraction process is then applied to perfectly subtract this signal at the receiver, before beamforming to the weaker transmission (avoiding the requirement to null the stronger, since it has already been subtracted). Finally the noisy QPSK data is detected at the output of this beamformer, which is a noisy received signal from the weaker transmit antenna.

Looking at the coded BLAST result in ~~Table 2~~Figure 17, by application of coding, it can be shown that the performance of BLAST has been improved by 4.9dB, down from 12.0dB to 7.1dB. This has been achieved by applying 'baseline' convolutional coding at the transmitter, along with a transfer of soft metrics from the beamformer pre-processor to the respective decoder. However, this is achieved at the expense of a reduction in spectral efficiency. Non-genie-aided simulation experience to date indicates that in practice the performance of coded BLAST 2:2 is in fact some 2-3dB worse than this, and so the subtraction processing falls rather short of a magic genie, due to the effect of error propagation down the layers. This would suggest that the transmission from the stronger transmit antenna (for any given frame) is actually more vulnerable (and therefore has a higher frame error probability) than the transmission from the weaker, due to the 'noise amplification' in the receiver antenna nulling process.

A13 Cont

~~Table 3~~Figure 18 shows a comparison of the complexities of the different STC algorithms in terms of MFLOPS (or MIPS if a fixed-point device is used). These processing loads assume a 153.6kbps data rate over the link. A 1-tap channel is assumed, and so the complexity analysis is pertinent to both TDMA and CDMA implementations in this flat channel. A number of simplifying assumptions have been made.

The processing tasks are partitioned into 'demodulator' -and 'Viterbi MLSE'. The demodulator generates the Log-Likelihood Ratios (LLR), also known as 'soft metrics', for the Viterbi decoder. The Viterbi algorithm then performs a maximum-

likelihood trellis search based on these LLRs in order to find the highest probability path through the trellis (the so-called 'survivor path'), and hence the most likely information sequence.

A13 cont'd

Examining the figures in ~~Table 3~~Figure 18, it is interesting to note that for the benchmark cases only a small percentage of the processing load is taken up by the demodulator and Maximal Ratio Combining (MRC), and the vast majority is in the Viterbi decoding (198.4 MFLOPS). Since the STTD 'de-rotation' pre-processing is similar in principle to MRC, this also doesn't add much to the overall processing. For all of these cases the Viterbi processing load is the same.

Page 28 lines 24-28:

A14

An STC scheme which is not ~~analysed~~analyzed in ~~Table 3~~Figure 18 is SC-STTC. However, experience with binary concatenated codes, such as Turbo codes, would indicate that the decoder complexity would certainly be well within an order of magnitude of the 'equivalent' non-Turbo (e.g. convolutional) code. The demodulator complexity (i.e. generation of soft metrics) for SC-STTC should therefore be of the same order as for the equivalent STTC.

Page 29 line 30 to page 30 line 8:

With respect to the implementation aspects of the invention, the following notes are relevant:

A15 cont

SAW filters: There is a requirement within STC for very tight control of delay through the multiple transmit chains. The main source of any such delay is believed to lie within the transmit SAW filters, since they are required to have steep roll-off skirts within their frequency response. SAW filters for these applications typically have a mean group delay (averaged over the pass-band) of the order of microseconds, for example. This mean group delay is much larger than the typical chip period for wideband CDMA applications. SAW filters also exhibit a considerable ripple in their frequency and phase (and hence group delay) response across their pass-band due

A15 cont'd

to triple-transit effects. However, it is believed that the wide-band nature of the CDMA waveform will mean that these delays are averaged out in the time domain, and that it is the mean delay which is of most concern, in particular how this can be matched across devices from the same, and possibly different, manufactured batches. SAW filter delay problems can be ~~minimised~~minimized by an appropriate choice of SAW filters which are well matched in terms of mean delay. An alternative could be to carry out some delay compensation in the transmitter DSP.